

# FAQ- Silicom Network Appliance



---

## *Silicom Network Appliance*

*featuring*

*A strong, yet standard commercial off the shelf host server mother board;*

*Coupled by 20 x 10GbE links to a capable switching silicon;*

*With up to 48 x 10GbE front I/O connectivity.*

---

FAQ

1. What version of software packages are supported? ..... 2
2. What about **NFV**? ..... 2
3. What are the advantages over using a standalone switch and standalone server?..... 3
4. Does it improve **VM to VM forwarding**? Also on separate CPU sockets? ..... 3
5. What about **Dual CPU server**? Isn't it enough? Isn't it as good as a switch?..... 3
6. What are your plans for next gen with RRC, what will change? ..... 4
7. Does it support **OpenFlow**? ..... 4
8. What about **SRIOV**? How many VMs are supported/accelerated by hardware? ..... 4
9. What about **MPLS**? Other tunneling? ..... 4
10. Is it a certified server with the common certification? ..... 4

1. What version of software packages are supported?

*Table 1 - Software Compatibility*

Software	Support	Version	Comment
<b><u>Operating Systems</u></b>			
Linux OS	✓	Ubuntu 14.04 Fedora 20	
<b><u>Hypervisors</u></b>			
KVM	✓		OVS offload option to hardware switch
XEN	✓		OVS offload option with XAPI
VMware	✓		No switching offload
<b><u>Network Engineering</u></b>			
OVS	✓		netdev interface to switching SDK
<b><u>Data Path</u></b>			
SPDK	✓	1.1	Silicom libraries for DPDK
DPDK	✓	1.7, 2.0	

2. What about **NFV**?

Silicom network appliance combines high powered server, high power switching fabric and large port density. All these enable it to support the three fundamentals of NFV:

- A) Virtual Network Function – As a COTS x86 server, most VNFs, whether for monitoring, SLA, security, etc. can adapt and run on Silicom network appliance.
- B) Close to the edge – Having switching silicon and a server in one chassis means being able to deploy virtual data center on this platform.
- C) Service Chaining – Same switching silicon is a key factor for efficient service chaining and packet brokerage.

### 3. What are the advantages over using a standalone switch and standalone server?

Managing a switching silicon on board immensely leverages network application traffic engineering capabilities, to its needs:

- A) **Traffic engineering:** Application may choose to use the capabilities of the switching silicon for elaborate traffic engineering to fit its needs. For instance, application may choose to sample traffic data, redirect ingress traffic, treat encapsulations (MPLS, VXLAN) differently, and much more.
- B) **Filtering:** The switching silicon features capable filtering engines, so that filtering tasks can be offloaded from host application to these engines.
- C) **Inter VM:** in a virtualized environment, a hypervisor may choose to offload all bridging and forwarding to a more capable switching silicon, designed just for that.
- D) **Simplified networking configuration:** Host holds persistent set of links to a switching entity.

### 4. Does it improve VM to VM forwarding? Also on separate CPU sockets?

Inter VM, or east-west traffic gradually increase in bandwidth. As more VMs interact internally within the virtualized environment, so increase the complexity of mesh switching and forwarding. This type of task involve table lookup and memory copy, that are all tasks to which a standard CPU is not best fitted. The benefit form offloading east-west forwarding to switching silicon, is twofold: First, CPU is relieved for processing tasks; and second, forwarding is performed by a best fitted piece of hardware.

### 5. What about Dual CPU server? Isn't it enough? Isn't it as good as a switch?

A standalone Dual CPU server has great capabilities indeed, and as a general purpose processing machine, it is capable of performing many types of tasks with the right software.

However, there are two major concerns involved when dealing with a Dual CPU server:

- A) **Quick Path Interconnect (QPI) bus:** Two CPU socket sharing the same mother board are interconnecting via QPI bus. Relatively speaking, this is a slow bus, and is used whenever CPU0 is required to process data residing on CPU1's RAM, and vice versa. In a scenario of busy processing or wire speed traffic, QPI bus inflicts its penalty.
- B) **CPU forté:** general purpose CPU is designed for linear processing. Simply that means that there are tasks to which is suits better than other tasks. The latters include parallel processing, and for our matter, network processing (table lookup, memory copy, etc.).

Therefore, great benefit is brought by the Silicom Network appliance design, and many things are achievable, that are beyond the capability of standalone CPU.

6. What are your plans for next gen with RRC, what will change?  
[TBD]

7. Does it support OpenFlow?

In one word, yes. This support is brought through several layers. First, there is the layer of the switching silicon, which is modern designed to support. Then, there is the layer the software API that operate this capability, and "speak" in flow syntax. Above it there comes the layer of OpenFlow protocol agents, such as OpenVswitch, that are accessing this API to configuration purposes. This all interoperability is demonstrated on Silicom Network Appliance.

8. What about SRIOV? How many VMs are supported/accelerated by hardware?

Five Intel® XL710 ("Fortville") network controllers implementing the network host interface in the NA26640, and as result, up to 32 SRIOV VFs per controller, or 160 total SRIOV VF in system.

9. What about MPLS? Other tunneling?

MPLS and VXLAN tunneling are identifiable by the filtering agent of the switching and network controllers, and so can be administrated and engineered. For instance, VXLAN tunnels can be forwarded and redirected according engineering policy, based on VNID.

10. Is it a certified server with the common certification?  
[TBD]